

Adaptive Horizon Model Predictive Control

Arthur J Krener

Abstract—Adaptive Horizon Model Predictive Control (AHMPC) is a scheme for varying as needed the horizon length of Model Predictive Control (MPC). Its goal is to achieve stabilization with horizons as small as possible so that MPC can be used on faster or more complicated dynamic processes. Beside the standard requirements of MPC including a terminal cost that is a control Lyapunov function, AHMPC requires a terminal feedback that turns the control Lyapunov function into a standard Lyapunov function in some domain around the operating point. But this domain need not be known explicitly. MPC does not compute off-line the optimal cost and the optimal feedback over a large domain instead it computes these quantities on-line when and where they are needed. AHMPC does not compute off-line the domain on which the terminal cost is a control Lyapunov function instead it computes on-line when a state is in this domain.

I. INTRODUCTION

Model Predictive Control (MPC) is a way to optimally steer a discrete time control system to a desired operating point. We briefly describe it following the definitive treatise of Rawlings and Mayne [1]. We closely follow their notation.

We are given a controlled, nonlinear dynamics in discrete time

$$x^+ = f(x, u) \quad (1)$$

where the state $x \in \mathbb{R}^{n \times 1}$, the control $u \in \mathbb{R}^{m \times 1}$ and $x^+(k) = x(k+1)$. This could be the discretization of a controlled, nonlinear dynamics in continuous time. The goal is to find a feedback law $u(k) = \kappa(x(k))$ that drives the state of the system to some desired operating point. A pair (x, u) is an operating point if $f(x, u) = x$. We conveniently assume that the operating point has been translated to be $(x, u) = (0, 0)$.

The controlled dynamics may be subject to constraints such as

$$x \in \mathbb{X} \subset \mathbb{R}^{n \times 1} \quad (2)$$

$$u \in \mathbb{U} \subset \mathbb{R}^{m \times 1} \quad (3)$$

and possibly a constraint involving both the state and control

$$y = h(x, u) \in \mathbb{Y} \subset \mathbb{R}^{p \times 1} \quad (4)$$

A control u is said to be feasible at $x \in \mathbb{X}$ if $u \in \mathbb{U}$ and

$$f(x, u) \in \mathbb{X}$$

$$h(x, u) \in \mathbb{Y}$$

Of course the stabilizing feedback $\kappa(x)$ that we seek needs to be feasible. For every $x \in \mathbb{X}$,

$$\begin{aligned} \kappa(x) &\in \mathbb{U} \\ f(x, \kappa(x)) &\in \mathbb{X} \\ h(x, \kappa(x)) &\in \mathbb{Y} \end{aligned}$$

An ideal way to solve this problem is to choose a Lagrangian $l(x, u)$ that is nonnegative definite in x, u and positive definite in u and then to solve the infinite time optimal control problem of minimizing over choice of feasible control sequence $\mathbf{u}_\infty = (u(0), u(1), \dots)$ the quantity

$$V_\infty(x) = \sum_{k=0}^{\infty} l(x(k), u(k)) \quad (5)$$

subject to the dynamics (1), the constraints (2, 4) and $x(0) = x$. Let $V_\infty^0(x)$ denote the minimum value and $\mathbf{u}_\infty^0 = (u_\infty^0(0), u_\infty^0(1), \dots)$ be a minimizing control sequence with corresponding state sequence $\mathbf{x}_\infty^0 = (x_\infty^0(0) = x, x_\infty^0(1), \dots)$. Minimizing control and state sequences need not be unique but we shall generally ignore this.

If a pair $V_\infty^0(x) \in \mathbb{R}, \kappa_\infty(x) \in \mathbb{R}^{m \times 1}$ of functions satisfy the infinite horizon Dynamic Program Equations (DPE_∞)

$$V_\infty^0(x) = \min_u \{l(x, u) + V_\infty^0(f(x, u))\} \quad (6)$$

$$\kappa_\infty(x) = \operatorname{argmin}_u \{l(x, u) + V_\infty^0(f(x, u))\} \quad (7)$$

$$V_\infty^0(0) = 0 \quad (8)$$

and the feasibility constraints

$$f(x, \kappa_\infty(x)) \in \mathbb{X} \quad (9)$$

$$h(x, \kappa_\infty(x)) \in \mathbb{Y} \quad (10)$$

for $x \in \mathbb{X}$ then it is not hard to show that $V_\infty^0(x)$ is the optimal cost and $\kappa_\infty(x)$ is an optimal feedback law, $u_\infty^0(k) = \kappa_\infty(x_\infty^0(k))$. Then under suitable conditions a Lyapunov argument can be used to show that the feedback $\kappa_\infty(x)$ is stabilizing.

The difficulty with this approach is that it is generally impossible to solve DPE_∞ on a large domain \mathbb{X} if the state dimension n is greater than 2 or 3. So both theorists and practitioners have turned to Model Predictive Control (MPC). They choose a Lagrangian $l(x, u)$, a horizon length N , a terminal domain $\mathbb{X}_f \subset \mathbb{X}$ containing $x = 0$ and a terminal cost $V_f(x)$ defined and positive definite on \mathbb{X}_f . Consider the problem of minimizing by choice of feasible $\mathbf{u}_N = (u_N(0), u_N(1), \dots, u_N(N-1))$

$$V_N(x) = \sum_{k=0}^{N-1} l(x(k), u(k)) + V_f(x(N)) \quad (11)$$

Research supported in part by AFOSR.

A. J. Krener is with the Department of Applied Mathematics, Naval Postgraduate School, Monterey, CA 93943 ajkrenner@nps.edu

subject to the dynamics (1), the constraints (2, 4), the terminal condition $x(N) \in \mathbb{X}_f$ and the initial condition $x(0) = x$. Assuming this problem is solvable, let $V_N^0(x)$ denote the optimal cost,

$$V_N^0(x) = \min_{\mathbf{u}_N} V_N(x) \quad (12)$$

where the minimum is taken over all feasible \mathbf{u}_N . Let $\mathbf{u}_N^0 = (u_N^0(0), u_N^0(1), \dots, u_N^0(N-1))$ and $\mathbf{x}_N^0 = (x_N^0(0) = x, x_N^0(1), \dots, x_N^0(N))$ denote optimal control and state sequences and define

$$\kappa_N(x) = u_N^0(0)$$

Let $\mathbf{X}_N \subset \mathbb{X}$ be defined inductively,

$$\mathbf{X}_0 = \mathbb{X}_f$$

$$\mathbf{X}_1 = \{x \in \mathbb{X} : \exists u \in \mathbb{U}, f(x, u) \in \mathbf{X}_0 \wedge (4)\}$$

$$\mathbf{X}_{N+1} = \{x \in \mathbb{X} : \exists u \in \mathbb{U}, f(x, u) \in \mathbf{X}_N \wedge (4)\}$$

The terminal set \mathbb{X}_f is controlled invariant (aka viable) if for each \mathbb{X}_f there exists a $u \in \mathbb{U}$ such that $f(x, u) \in \mathbb{X}_f$ and the constraints (4) are satisfied. If this holds then it is not hard to see inductively that the sets are nested $\mathbf{X}_N \subset \mathbf{X}_{N+1}$.

If a pair $V_N^0(x), \kappa_N(x)$ defined on \mathbf{X}_N satisfy the horizon N Dynamic Program Equations (DPE_N)

$$V_N^0(x) = \min_u \{l(x, u) + V_N^0(f(x, u))\} \quad (13)$$

$$\kappa_N(x) = \operatorname{argmin}_u \{l(x, u) + V_N^0(f(x, u))\} \quad (14)$$

$$V_N^0(x) = V_f(x) \text{ for } x \in \mathbb{X}_f \quad (15)$$

where the minimum is over all $u \in \mathbb{U}$ that are feasible at $x \in \mathbf{X}_N$ then it is not hard to show that $V_N^0(x)$ is the optimal cost and $\kappa_N(x)$ is an optimal feedback law $u_N^0(k) = \kappa(x_N^0(k))$. If $V_f(x)$ is a control Lyapunov function on \mathbb{X}_f then under suitable conditions a Lyapunov argument can be used to show that the feedback $\kappa_N(x)$ is stabilizing on \mathbf{X}_N . See [1] for more details.

As we noted above solving off-line the infinite horizon optimal control problem for all possible states is generally intractable. The advantage of solving the horizon N optimal control problem for the current state x is that it possibly can be done on-line as the process evolves. If the current value of the state is known to be $x \in \mathbf{X}_N$ then the finite horizon N optimal control problem is a nonlinear program with finite dimensional decision variable $\mathbf{u}_N \in \mathbb{R}^{m \times N}$. If the time step is long enough, if f, h, l are reasonably simple and if N is small enough then this nonlinear program that can be solved in a fraction of one time step for \mathbf{u}_N^0 . Then the first element of this sequence $u_N^0(0)$ is used as the control at the current time. The system evolves one time step and the process is repeated at the next time. Conceptually MPC computes an optimal feedback law $\kappa_N(x) = u_N^0(0)$ but only at values of x when and where it is needed.

Some authors do away with the terminal cost $V_f(x)$ but there is a theoretical and a practical reason to use one. The theoretical reason is that a control Lyapunov terminal cost facilitates a proof of asymptotic stability via a Lyapunov argument [1]. The practical reason is that one can usually

use a shorter horizon N when there is a terminal cost. A shorter horizon reduces the dimension mN of the decision variables in the nonlinear programs that need to be solved on-line. Therefore MPC with a suitable terminal cost can be used for faster and more complicated systems.

The ideal terminal cost $V_f(x)$ is $V_\infty(x)$ of the corresponding infinite horizon optimal control provided that the latter can be accurately computed off-line on a reasonably large terminal set \mathbb{X}_f . This may be tractable because the terminal set \mathbb{X}_f may be much smaller than \mathbb{X} and only an approximate solution on \mathbb{X}_f may suffice. For example $V_\infty(x)$ can be locally approximated by the solution of the infinite horizon LQR problem involving the linear part of the dynamics and quadratic part of the Lagrangian at the operating point.

One would expect when the current state x is far from the operating point, a relatively long horizon N is needed to ensure that $x_N^0(N) \in \mathbb{X}_f$ but as the state approaches the operating point shorter and shorter horizons can be used. Adaptive Horizon Model Predictive Control (AHMPC) adjusts the horizon of MPC on-line as it is needed. In the next section we present an ideal version of AHMPC and in the following section we present a practical implementation of AHMPC. Finally we close with an example.

II. IDEAL ADAPTIVE MODEL PREDICTION CONTROL

We shall make some standing assumptions. The first few are drawn from Rawlings and Mayne.

Assumption 1: (Assumption 2.2 [1])

The functions $f(x, u), l(x, u), h(x, u), V_f(x)$ are continuous on some open set containing $\mathbb{X} \times \mathbb{U}$, $l(x, u)$ is nonnegative definite in (x, u) and positive definite in u on this open set, V_f is positive definite on \mathbb{X}_f and $f(0, 0) = 0, l(0, 0) = 0, V_f(0) = 0$.

Assumption 2: (Assumption 2.3 [1])

The sets \mathbb{X} and \mathbb{X}_f are closed, $\mathbb{X}_f \subset \mathbb{X}$, \mathbb{U} is compact and \mathbb{X}, \mathbb{X}_f and \mathbb{U} contain neighborhoods of their respective origins.

Assumption 3: (Assumptions 2.12 and 2.13 of [1])

For all $x \in \mathbb{X}_f$ there exist a feasible u such that

$$\begin{aligned} f(x, u) &\in \mathbb{X}_f \\ l(x, u) + V_f(f(x, u)) &\leq V_f(x) \end{aligned}$$

This assumption implies that \mathbb{X}_f is controlled invariant and that $V_f(x)$ is a control Lyapunov function on \mathbb{X}_f .

We make some additional assumptions.

Assumption 4: For each $x \in \mathbb{X}$ there is a nonnegative integer N and a control sequence $\mathbf{u}_N = (u_N(0), \dots, u_N(N-1))$ such that the corresponding state sequence $\mathbf{x}_N = (x_N(0) = x, \dots, x_N(N))$ starting from x satisfies $x_N(N) \in \mathbb{X}_f$.

This assumption allows us to define a function $N(x)$ on \mathbb{X} as the minimum of all N such that there exist such a control sequence $\mathbf{u}_N = (u_N(0), \dots, u_N(N-1))$ and corresponding state sequence $\mathbf{x}_N = (x_N(0) = x, \dots, x_N(N))$ starting from x that satisfies $x_N(N) \in \mathbb{X}_f$.

Then the nested sets \mathbf{X}_N defined above are given by

$$\mathbf{X}_N = \{x \in \mathbb{X} : N(x) \leq N\}$$

Assumption 5: There exists a nonnegative integer M such that $N(x) \leq M$ for all $x \in \mathbb{X}$. In other words

$$\mathbf{X}_M = \mathbb{X}$$

These assumptions imply that the usual MPC with horizon length M is stabilizing on \mathbb{X} by standard arguments [1]. But it is a waste of time to use horizon length M when $N(x)$ is substantially smaller M . If the current state is x , ideal AHMPC uses horizon length $N(x)$. Then as the current state approaches the terminal set, ideal AHMPC uses shorter and shorter horizons. When x is in the terminal set, ideal AHMPC uses a horizon length of $N = 0$.

In a moment we shall show that the function

$$V(x) = V_{N(x)}^0(x) \quad (16)$$

is a valid Lyapunov function for the closed loop system which confirms the stabilizing property of the ideal AHMPC feedback

$$\kappa(x) = \kappa_{N(x)}(x) \quad (17)$$

But the reason why this scheme is not practical is that, in general, it is impossible to compute the function $N(x)$. In the next section we shall offer a work around but for now we study the stabilizing properties of ideal AHMPC.

Lemma 1 Assume Assumptions 1-5 hold. If $N(x) = N$ and if $\mathbf{u}_N = (u_N(0), \dots, u_N(N-1))$, $\mathbf{x}_N = (x_N(0) = x, \dots, x_N(N))$ are a control and state trajectory from x such that $x_N(N) \in \mathbb{X}$ then $N(x(k)) = N - k$.

Proof: By assumption $x_N^0(N) \in \mathbb{X}_f$ so $x_N^0(N-1) \in \mathbf{X}_1$, $x_N^0(N-2) \in \mathbf{X}_2$, etc. So $N(x_N^0(k)) \leq N - k$.

Suppose for some k , $x_N^0(k) < N - k$ then $x_N^0(k) \in X_{N-k-1}$ and $x_N^0(k-1) \in X_{N-k-2}$, $x_N^0(k-2) \in X_{N-k-3}$, etc. so $x = x_N^0(0) \in X_{N-1}$ which contradicts the assumption that $N(x) = N$. ■

Lemma 2 (Compare with Lemma 2.14 of [1]) Under Assumptions 1-5 then

$$V(x^+) \leq V(x) - l(x, \kappa(x))$$

where $x^+ = f(x, \kappa(x))$ and V is defined by (16).

Proof: By definition

$$\begin{aligned} V(x) &= V_{N(x)}^0(x) \\ &= \sum_{k=0}^{N(x)-1} l(x_{N(x)}^0(k), u_{N(x)}^0(k)) + V_f(x_{N(x)}^0(N(x))) \end{aligned}$$

where $\mathbf{u}_{N(x)}^0 = (u_{N(x)}^0(0), \dots, u_{N(x)}^0(N(x) - 1))$ and $\mathbf{x}_{N(x)} = (x_{N(x)}(0) = x, \dots, x_{N(x)}(N(x)))$ are optimizing control and state sequences for the horizon $N(x)$ optimal control problem so $\kappa(x) = u_{N(x)}^0(0)$ where κ is defined by (17).

Let $x^+ = f(x, \kappa(x)) = x_{N(x)}^0(1)$, by Lemma 1, $N(x^+) = N(x) - 1$ so

$$V(x^+) = V_{N(x)-1}^0(x^+)$$

So for any feasible control sequence $\bar{\mathbf{u}} = (u(0), \dots, u(N(x) - 2))$ and corresponding state sequence $\bar{\mathbf{x}} = (x(0) = x^+, \dots, x(N(x) - 1))$

$$V(x^+) \leq \sum_{k=0}^{N(x^+)-1} l(x(k), u(k)) + V_f(x_{N(x)}(N(x)))$$

In particular if we take $\bar{\mathbf{u}} = (u_{N(x)}^0(1), \dots, u_{N(x)}^0(N(x) - 1))$ then $\bar{\mathbf{x}} = (x_{N(x)}(1), \dots, x_{N(x)}(N(x)))$ and

$$V(x^+) \leq V(x) - l(x, \kappa(x)) \quad \blacksquare$$

Following Rawlings and Mayne we make the following assumption.

Assumption 6: (Assumption 2.16(a) of [1]) The stage cost l and the terminal cost V_f satisfy

$$\begin{aligned} l(x, u) &\geq \alpha_1(|x|) \quad \forall x \in \mathbb{X}, \forall u \in \mathbb{U} \\ V_f(x) &\leq \alpha_2(|x|) \quad \forall x \in \mathbb{X}_f \end{aligned}$$

where $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$ are K_∞ functions.

Assumptions 3 and 6 imply that for each $x \in \mathbb{X}_f$ there exists a feasible u such that

$$V(f(x, u)) \leq V(x) - \alpha_1(|x|)$$

Proposition 1 (Compare with Proposition 2.17 of [1]) (a) Suppose that Assumptions 1, 2, 3, 4, 5 and 6 are satisfied. Then there exists K_∞ functions $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$ such that $V(\cdot)$ has the following properties

$$\begin{aligned} V(x) &\geq \alpha_1(|x|) \quad \forall x \in \mathbb{X} \\ V(x) &\leq \alpha_2(|x|) \quad \forall x \in \mathbb{X}_f \\ V(f(x, \kappa(x))) &\leq V(x) - \alpha_1(|x|) \quad \forall x \in \mathbb{X} \end{aligned}$$

Proof If $N(x) > 0$ then the first inequality follows from Assumption 6(a) and the fact that $V(x) \geq l(x, \kappa(x))$. If $N(x) = 0$ then the first inequality follows from Assumption 7(a). The second inequality follows Assumption 6(a) and the fact that if $x \in \mathbb{X}_f$ then $N(x) = 0$ and so $V(x) = V_f(x)$. ■

If the second property held for all $x \in \mathbb{X}$

$$V(x) \leq \alpha_2(|x|) \quad \forall x \in \mathbb{X}$$

then $V(x)$ would be a valid Lyapunov on \mathbb{X} . The following is a paraphrase of a proposition of Rawlings and Mayne.

Proposition 2 (Proposition 2.1 of [1])

Suppose that Assumptions 1, 2, 3 hold, that \mathbb{X}_f contains an open neighborhood of the origin and that \mathbb{X} is compact. If there exists a K_∞ function $\alpha(\cdot)$ such that $V_N^0(x) \leq \alpha(|x|)$ for $x \in \mathbb{X}_f$ then there exists another K_∞ function $\beta_N(\cdot)$ such that $V_N^0(x) \leq \beta_N(|x|)$ for $x \in \mathbf{X}_N$.

This allows us to prove the following proposition.

Proposition 3

Suppose that Assumptions 1, 2, 3, 4, 5 and 6 hold, that \mathbb{X}_f contains an open neighborhood of the origin and that

\mathbb{X} is compact. If there exists a K_∞ function $\alpha(\cdot)$ such that $V_N^0(x) \leq \alpha(|x|)$ for $x \in \mathbb{X}_f$ then there exists another K_∞ function $\beta(\cdot)$ such that $V_N^0(x) \leq \beta(|x|)$ for $x \in \mathbb{X}$.

Proof: By Assumptions 4 and 5, $\mathbb{X} = \mathbf{X}_M$. Let $\beta_0(\cdot) = \alpha(\cdot)$ and define

$$\beta(s) = \max \{\beta_N(s) : N = 0, 1, \dots, M\}$$

The maximum of a finite family of K_∞ functions is also a K_∞ function. Clearly if $x \in \mathbb{X} = \mathbf{X}_M$

$$V(x) = V_{N(x)}^0(x) \leq \beta_{N(x)}(|x|) \leq \beta(|x|) \quad \blacksquare$$

Proposition 4 Suppose that Assumptions 1, 2, 3, 4, 5 and 6 hold, that \mathbb{X}_f contains an open neighborhood of the origin and that \mathbb{X} is compact. Then $V(x)$ is a valid Lyapunov function which confirms the asymptotic stability of the closed loop dynamics

$$x^+ = f(x, \kappa(x))$$

on \mathbb{X}

So ideal AHMPC solves our stabilization problem in theory. But it generally can't be implemented because we can't compute the key ingredient, the function $N(x)$ or its domain of definition.

There is a slightly less ideal version of AHMPC. Suppose we have a function $N(x)$ with the following properties.

- For each $x \in \mathbb{X}$ there is a feasible control sequence $\mathbf{u}_{N(x)} = (u_{N(x)}(0), \dots, u_{N(x)}(N(x) - 1))$ and corresponding state sequence $\mathbf{x}_{N(x)} = (x_{N(x)}(0) = x, \dots, x_{N(x)}(N))$ such that $x_{N(x)}(N(x)) \in \mathbb{X}_f$.
- There exist an M such that $N(x) \leq M$ for all $x \in \mathbb{X}$.
- If $\mathbf{u}_{N(x)}^0 = (u_{N(x)}^0(0), \dots, u_{N(x)}^0(N(x) - 1))$ and $\mathbf{x}_{N(x)}^0 = (x_{N(x)}^0(0) = x, \dots, x_{N(x)}^0(N))$ are optimal control and state sequences for the horizon $N(x)$ optimal control problem starting at x then $0 \leq N(x_{N(x)}(k)) - N(x_{N(x)}(k + 1)) \leq 1$. In other words along optimal trajectories $N(\cdot)$ either stays the same or decreases by 1 at each time step.

Then the above results hold, $V(x)$ as defined by (16) is a valid Lyapunov function for closed loop system using feedback $\kappa(x)$ defined by (17). The only additional thing that needs to shown is that if $N(x) = N(x^+)$ then $V(x) - V(x^+) \geq \alpha_1(|x|)$. But this follows from standard MPC arguments, see Lemma 2.14 of [1].

III. ADAPTIVE HORIZON MODEL PREDICTIVE CONTROL

Here is a variation on the above that is practical which we call Adaptive Horizon Model Predictive Control (AHMPC). We assume that we have the following.

- Sets \mathbb{X} , \mathbb{X}_f , \mathbb{U} satisfying Assumption 2. We do not require that \mathbb{X}_f be known explicitly.
- A discrete time controlled dynamics $f(x, u)$, a Lagrangian $l(x, u)$, a constraint pair $(h(x, u), \mathbb{Y})$ and a terminal cost $V_f(x)$ satisfying Assumption 1.

- A terminal feedback $u = \kappa_f(x)$ and a class K_∞ function $\alpha(\cdot)$ defined for all $x \in \mathbb{X}_f$ and satisfying

$$\begin{aligned} V_f(x) &\geq \alpha(|x|) \\ f(x, \kappa_f(x)) &\in \mathbb{X}_f \\ V_f(x) - V_f(f(x, \kappa_f(x))) &\geq \alpha(|x|) \\ h(x, \kappa_f(x)) &\in \mathbb{Y} \end{aligned}$$

We don't need to know the terminal set \mathbb{X}_f on which these conditions are satisfied, all we need to there is such a terminal set and that it contains a neighborhood of $x = 0$.

One way of obtaining such a terminal pair $V_f(x)$, $\kappa_f(x)$ is to approximately solve the infinite horizon dynamic program equations (DPE $_\infty$) on some neighborhood of the origin. For example if the linear part of the dynamics and the quadratic part of the Lagrangian constitute a nice LQR problem then then one can let $V_f(x)$ be the quadratic optimal cost and $\kappa_f(x)$ be the linear optimal feedback of the LQR. Alternatively one can take higher degree Al'brenkt approximations to $V_\infty(x)$, $\kappa_\infty(x)$ [2]. Of course the problem with such terminal pairs $V_f(x)$, $\kappa_f(x)$ is that generally there is no way to estimate the terminal set \mathbb{X}_f on which (1), (2) and (3) are satisfied. It is reasonable to expect that they are satisfied on some terminal set but the extent of the terminal set is very difficult to estimate.

AHMPC mitigates this difficulty. MPC does not try to compute the optimal cost and optimal feedback everywhere, instead it computes them just when and where they are needed. AHMPC does not try to compute the extent of V_f , it just tries to determine if the end state $x_N^0(N)$ of the currently computed optimal trajectory is in a terminal set \mathbb{X}_f where (1), (2) and (3) are satisfied.

Suppose the current state is x and we have solved the horizon N optimal control problem for $\mathbf{u}_N^0 = (u_N^0(0), \dots, u_N^0(N - 1))$, $\mathbf{x}_N^0 = (x_N^0(0) = x, \dots, x_N^0(N))$. AHMPC does not explicitly impose the terminal constraint $x_N^0(N) \in \mathbb{X}_f$ because \mathbb{X}_f is not explicitly known but it does require that the terminal cost V_f is defined at $x_N^0(N)$.

The terminal feedback $u = \kappa_f(x)$ is used to extend the state trajectory L additional steps

$$x_N^0(k + 1) = f(x_N^0(k), \kappa_f(x_N^0(k)))$$

for $k = N, \dots, N + L - 1$. This assumes that the terminal feedback is defined $u = \kappa_f(x)$ is defined on $x_N^0(k)$ for $k = N, \dots, N + L - 1$. If the terminal feedback is not defined at any of these points then we presume that $x_N^0(N)$ is not in \mathbb{X}_f so we increase N by 1 and we solve the optimal control problem over the new horizon.

If the feedback is defined on the extended trajectory then one checks that the Lyapunov conditions hold for the extended part of the state sequence,

$$V_f(x_N^0(k)) \geq \alpha(|x_N^0(k)|) \quad (18)$$

$$V_f(x_N^0(k)) - V_f(x_N^0(k + 1)) \geq \alpha(|x_N^0(k)|) \quad (19)$$

for $k = N, \dots, N + L - 1$. Again if the terminal cost V_f is not defined at any of these points then we presume that

$x_N^0(N)$ is not in \mathbb{X}_f so we increase N by 1 and we solve the optimal control problem over the new horizon.

If (18, 19) hold for all for $k = N, \dots, N + L - 1$. then we presume that $x_N^0(N) \in \mathbb{X}_f$ and we use the control $u_N^0(0)$ to move one time step forward to $x^+ = f(x, u_N^0(0))$. At this next state x^+ we solve the horizon $N - 1$ optimal control problem and check that the extension of the new optimal trajectory satisfies (18, 19).

If (18, 19) do not hold for all for $k = N, \dots, N + L - 1$. then we presume that $x_N^0(N) \notin \mathbb{X}_f$. If time permits we solve the horizon $N + 1$ optimal control problem at the current state x and then check the Lyapunov conditions (18, 19) again. We keep increasing the horizon by 1 until these conditions are satisfied. If we run out of time before (18, 19) are satisfied then we use the last computed $u_N^0(0)$ and move one time step forward to $x^+ = f(x, u_N^0(0))$. At x^+ we solve the horizon $N + 1$ optimal control problem.

The number L of additional time steps is a design parameter. Two obvious choices are to take a fixed L which is a fraction of M or to take a varying L which is a fraction of the current N .

IV. EXAMPLE

The example that we apply AHMPC to is stabilizing a double pendulum to the upright position using torques at each of the pivots. The states are x_1 , the angle of the first leg measured in radians counter-clockwise from straight up, x_2 , the angle of the second leg measured in radians counter-clockwise from straight up, $x_3 = \dot{x}_1$ and $x_4 = \dot{x}_2$. The controls are u_1 , the torque applied at the base of the first leg, and u_2 , the torque applied at the joint between the legs. The length of the first leg is 1 m. and the length of the second leg is 2 m. The legs are assumed to be massless but there is a mass of 2 kg. at the joint between the legs and a mass of 1 kg. at the tip of the second leg. The continuous time controlled dynamics is discretized using Euler's method with time step 0.1 s. assuming the control is constant throughout the time step.

The continuous time Lagrangian is chosen to be $l_c(x, u) = (|x|^2 + |u|^2)/2$ and its Euler discretization, $l(x, u) = (|x|^2 + |u|^2)/20$, is used. We choose the initial state to be $x = (\pi/2, -\pi/2, 0, 0)'$ and the initial horizon length to be $N = 5$. We simulated practical AHMPC with $V_f(x)$, $\kappa_f(x)$ the solution of the LQR problem using the linear part of the dynamics at the origin and the quadratic Lagrangian, $\alpha(|x|) = 0.1|x|^2$ and fixed $L = 5$. We did not move one time step forward if (18, 19) did not hold over the extended state trajectory but instead increased N by one and recomputed. The AHMPC trajectories of the two angles, x_1 in blue and x_2 in red, are shown in Figure 1. The adaptively changing horizon length is shown in Figure 2. This graph includes cases where the horizon was increased by one but the state of the pendulum was not advanced. Notice that the horizon goes down and up several times before settling at $N = 0$.




Fig. 1. Angles Converging to the Vertical




Fig. 2. Adaptively Changing Horizon

V. CONCLUSION

Adaptive Horizon Model Predictive Control is a scheme for varying the horizon length in Model Predictive Control as the stabilization process evolves. We have presented an ideal version of AHMPC and shown that it guarantees stabilization. AHMPC is a practical version that proceeds without knowing the minimum horizon length function $N(x)$ and without knowing the domain of Lyapunov stability of the terminal cost $V_f(x)$ and terminal feedback $\kappa_f(x)$.

We have only proven the convergence of AHMPC under ideal conditions but the convergence of standard MPC is also proven under similar ideal conditions, e.g., exact model, exact knowledge of the current state, exact solution of the finite horizon optimal control problems, etc.

The principal advantage of AHMPC over standard MPC is that the AHMPC horizon length decreases as the process is stabilized thereby lessening the on-line computational

burden. Hence AHMPC may be able to stabilize systems with faster or more complicated dynamics.

The author would like to acknowledge helpful communications with Sergio Lucia, Philipp Rumschinski and Rolf Findeisen.

REFERENCES

- [1] J. B. Rawlings and D. Q. Mayne, Model predictive control : theory and design, Nob Hill Pub., 2009.
- [2] E. G. Al'brekht, On the Optimal Stabilization of Nonlinear Systems, PMM-J. Appl. Math. Mech., 25:1254-1266, 1961.